Research Subcontracts

The ISCR supported 19 research subcontracts to various institutions throughout the United States. These contracts are normally funded by programs at LLNL to help address long-term Laboratory requirements. These subcontracts typically fund residential visits by university faculty for close collaborations with scientists in the Computation Directorate. Brief reports follow detailing ongoing work enabled by these subcontracts.

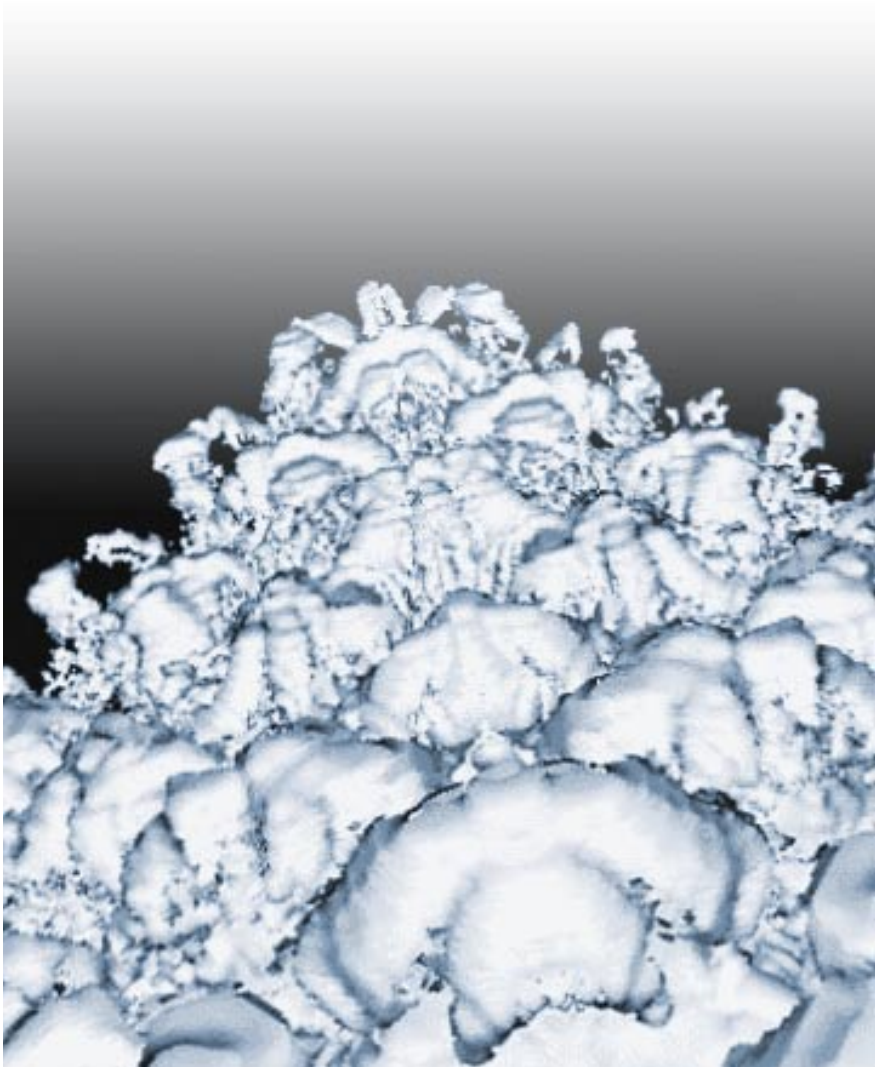

Nanohydrodynamic Simulation of the Rayleigh–Taylor Instability

Principal Investigator

Berni Alder, University of California, Davis

Collaborator

Garry Rodrigue, CASC

During the past year, we collaborated with a group at Los Alamos National Laboratory (LANL) and successfully carried out molecular dynamics simulations of the Rayleigh–Taylor instability that quantitatively predicts both the early-time behavior given by linear stability analyses and the long-time behavior given by experiment, as well as some of the numerical solutions of the continuum Navier–Stokes equation. This was published in *Proceedings of the National Academy of Sciences* and helped establish why different long-time behavior was observed in various continuum calculations.

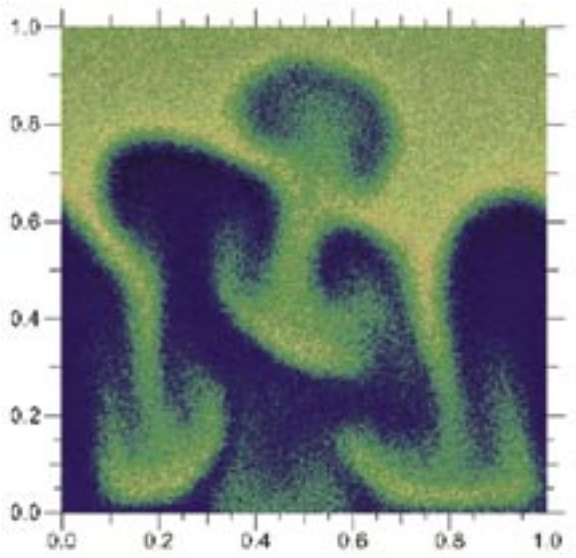

Demonstration of the Rayleigh–Taylor instability. The color variation goes from yellow at the top for the highest density to dark blue for the lowest density.

This is the first time that three-dimensional particle methods quantitatively simulated turbulent mixing. In the biggest and only molecular dynamics run, a qualitatively different asymptotic behavior was found, namely the heavy fluid formed drops instead of coalescing into a single slug. This could be real and appear only in particle simulations due to the presence of fluctuations, which could cut the thin thread by which the mushrooms are connected to the tip of the spike.

In order to firmly establish this phenomenon, which is of great importance to groups such as the NIF and because the mixing region advances more slowly (i.e., time to the first instead of the second power), we have developed an alternative particle algorithm that is 100 times faster and allows us to study even bigger systems for longer times. This was accepted for publication in *Molecular Physics*. To show the power of this scheme, we have attached a graph of the Rayleigh–Taylor instability at fairly long times with a few million particles that were produced on a single–processor computer in a few hours.

What these preliminary results indicate is the necessity to confirm them by carrying out the biggest Rayleigh–Taylor calculation possible on the IBM BlueGene/L machine for which the new particle algorithm is particularly well suited. We also need to carry out a comparison to the best continuum solution to see the quantitative effects of the nonlinear terms and fluctuations left out in the Navier–Stokes equations.

Multiple Animal Tracking in the Smart Vivarium

Principal Investigator

Serge Belongie, University of California, San Diego

Collaborator

Chandrika Kamath, CASC

A common trend in 2D object recognition is to detect and leverage the use of sparse, informative feature points. The use of such features makes the problem more manageable while providing increased robustness to noise and pose variation. In this work, we extend these ideas to the spatiotemporal case. For this purpose, we show that the direct 3D counterparts to commonly used 2D interest-point detectors are inadequate and propose an alternative. Anchoring off of these interest points, we devise a recognition algorithm based on spatio-temporally windowed data. We present promising recognition results on a challenging real-world database of mouse behaviors.

We attempt to extend the above approaches developed for object recognition to the problem of behavior recognition, i.e., from the spatial to the spatiotemporal domain. These extensions are not always direct, but rather follow the general spirit of using sparsely detected features for object recognition. We propose to characterize behavior through the use of spatiotemporal feature points.

For the above purpose, we show that the direct 3D counterparts to commonly used 2D interest-point detectors are inadequate and propose an alternative. We extend descriptors for spatial interest points to cuboids of spatiotemporally windowed data. Cuboids extracted from a number of sample behaviors are clustered and the resulting cluster centers serve as a dictionary of atomic units of behavior we call "actons." The only information kept from all subsequent video data is the location and type of the actons present. We argue that such a representation is sufficient for recognition and robust with respect to variations in the data. We show an application of this framework, utilizing a simple classification scheme, to a challenging real-world database of mouse behaviors.

Given these feature vectors and labelled training data, we trained two different classifiers: linear discriminant analysis and a DAG support vector machine using a radial basis kernel. Both were applied after dimensionality reduction using principal components analysis. As mentioned in each case,six clips were used for training and the seventh for testing. We limited the amount of training data for the over-represented categories, such as exploring, otherwise categories such as drinking would be drowned out. However, we used all the testing data. The SVM classifier outperformed LDA by a few percentage points. The classifiers achieved an accuracy ranging from about 75% to 95%.

Robust Trajectory and Appearance-Based Data Association for Multi-Object Tracking

Principal Investigator **Michael C. Burl,** University of Colorado, Boulder

Collaborator **Chandrika Kamath,** CASC

Reliable tracking is clearly a prerequisite for more advanced video-mining operations. This study builds upon our previous work in multi-object tracking from surveillance-style video sequences. The primary goal was to improve upon the tracker's ability to assign and maintain an object's unique track ID for the duration of time that the object appears in the scene.

The general structure of our previous tracker consisted of six steps.

1. Background estimation and subtraction
2. Thresholding and spatial grouping of detected pixels into blobs
3. Gating and data association to link blobs with tracks
4. Birth of new tracks if there are unassociated blobs
5. Measurement update (incorporating new observations into the state estimates for each of the tracks)
6. Time update (predicting the position and appearance of the tracked objects at the next observation time).

A key problem with this type of tracker is that the blobbing process is not perfect and leads to over-segmentation (a single object split into multiple blobs) and under-segmentation (multiple objects lumped into a single blob). Subsequent steps in our previous tracker could not properly deal with these fundamental grouping errors.

In this study, we investigated a number of ideas that we believed could improve tracking performance. A two-pronged approach was pursued with part of the effort devoted to getting more information from an object's trajectory and part of the effort devoted toward getting more information from an object's appearance.

The following component ideas were considered.

- Compensating for perspective effects
- Using more refined dynamical models
- Using decision-theoretic local search (structured change detection)
- Establishing precise correspondence between parts of a tracked object from one frame to the next, splitting and merging blobs
- Using particle filters to overcome some of the limitations caused by hard data association decisions.

Although we explored these ideas in some detail, we were not able to integrate these disparate components together into a unified tracker within the allotted time, nor were we able to conduct systematic evaluations over benchmark data sets to characterize the overall impact on performance. However, we plan to continue working toward these two goals.

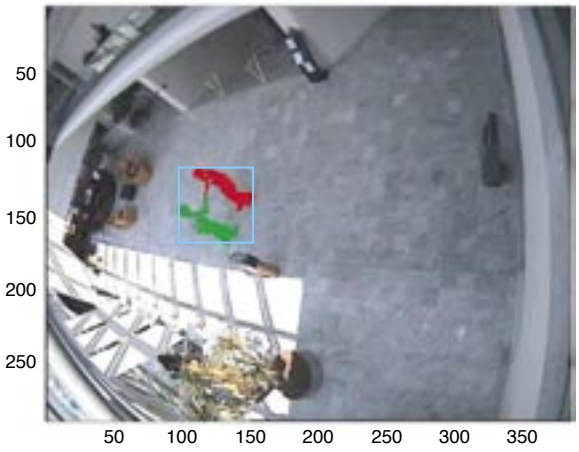

Two people meet and shake hands. The object detection component initially groups the two people together into a single blob (cyan bounding box), but the tracker is able to use the appearance and trajectory information from the previous frames to correctly split the blob into two individuals.

Multi-Constraint and Multi-Objective Partitioning-Complex Networks

Principal Investigator

Umit Catalyurek, Ohio State University

Collaborator

Edmond Chow, CASC

The aim of this project was to get better insight into partitioning properties of complex networks. With the advances in high-performance computing and high-resolution sensors in many fields of the computational sciences, data sizes are getting bigger and bigger. The sizes of complex networks that arise in homeland security applications are no different than any the other applications. Both the size and on-demand query and analysis requirements of applications that use complex networks necessitate both parallel and distributed computing. To achieve good performance, data needs to be partitioned in a load-balancing manner and communication requirements needs to be minimized.

Starting with random networks, we investigated the performance of 1D and 2D coarse-grain hypergraph partitioning of various networks provided to us by LLNL. Using the Performing Tools for Hypergraph (PaToH) partitioning tool, we analyzed the change in two cost metrics—graph edge cut and hypergraph connectivity-1—with respect to several graph parameters and the number of partitions. We also derived theoretical relationships between those parameters and the cost metrics.

In the experimentation phase for the type of the networks that have a generator, we generated multiple graphs with different parameter values, such as average degree and number of vertices. For each graph, we first used a random partitioner to generate a partition that would be used as a baseline to evaluate the performance of PaToH. Next, we partitioned these graphs using 1D and 2D coarse-grain hypergraph models using PaToH as our partitioner. Finally, we compared the performances of those partitioning schemes and explained the results in accordance with the theoretical predictions.

Our results showed that both 1D and 2D partitioning achieve much better results than the random partitioning, even though both approaches produce similar communication volume requirements. However, in applications that are sensitive to the number of communications, partitioning via 2D coarse-grain decomposition achieves much better results, providing better performance. Searching through paths in complex networks, especially for the homeland security applications that motivate our work, provides a practical real-world scenario where partitioning using 2D coarse-grain decomposition would be advantageous.

Numerical Simulation of Dispersion of Solid Particles Moving at Supersonic Speeds in Turbulent Flows

Principal Investigator **Said Elghobashi,** University of California, Irvine

Collaborator **Bill Bateson,** CASC

This work aims at developing a mathematical model that predicts the trajectories of spherical solid particles moving at supersonic speeds in a turbulent flow. The mathematical model involves the numerical solution of the three-dimensional, time-dependent, compressible, Navier–Stokes and energy equations. The project comprises two main tasks to be performed in parallel. Each task will require four years for completion.

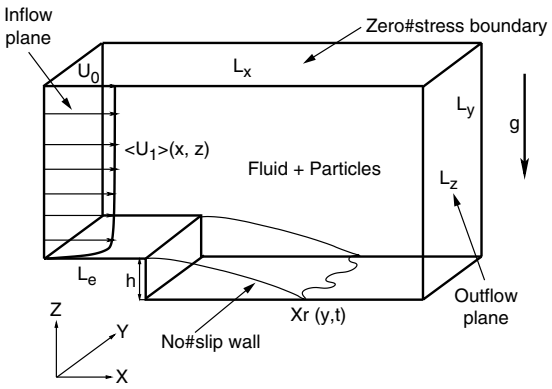

Figure 1. Schematic of the computational domain.

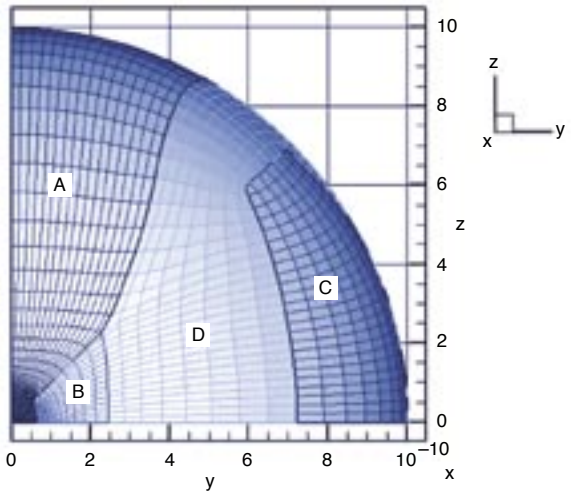

Figure 2. Grid surface ZY plane view

Task 1. The objective of this task is to compute the unsteady 3D compressible flow around a single-fixed sphere subjected to an air flow of a speed about 5000 m/s. An existing code at University of California, Irvine has been modified for a laminar compressible flow to account for the presence of a fixed sphere. Integrating the forces imparted on the sphere surface provides the values of the drag and lift on the sphere. The effects of varying the Mach number and particle Reynolds number will be investigated. A differential equation of motion of a sphere at supersonic speeds ($M > 1$) will be developed based on the computed drag and lift forces.

Task 2. The objective of this task is to compute the 3D turbulent isotropic compressible flow laden with many spherical solid particles moving at supersonic speeds in a cubical domain. The motion of each particle will be governed by the differential equation mentioned above in Task 1. A mathematical model for the collision between the particles will be included in the particle motion equation.

The detailed report describes the mathematical approach and the results of an incompressible flow over a single sphere. The computed values of the drag coefficient agree well with those of Kim, et.al. for the range of Reynolds numbers Re = 20 to 100.

Analysis and Visualization of Scientific Data Sets Using Generalized Segmentation Methods

Principal Investigator **Bernd Hamann,** University of California, Davis

Collaborator **Valerio Pascucci,** CASC

Segmentation of scientific data sets is becoming increasingly important in the context of understanding massive numerically simulated data sets at higher levels of qualitative behavior. Topology is one area in mathematics that allows us to characterize the behavior of scalar-valued data sets based on analyzing critical points and their relationships. This effort developed new techniques and program modules that support the automatic and mathematically sound segmentation of scalar-valued data sets into subsets (subregions of the physical domain) of topologically (and thus qualitatively) distinct behavior.

Traditional volume rendering of a data set involves visualizing surface properties or utilizing a transfer function. An alternative approach visualizes topological information, supporting a more qualitative understanding. The Morse–Smale complex is a structure that represents the topology of a scalar data set. The Morse–Smale complex can be simplified for use in applications, such as topological smoothing or hierarchical data visualization.

We have developed algorithms for extracting and simplifying Morse–Smale complexes that rely solely on combinatorial decisions, therefore avoiding numerical instability. We extract a Morse–Smale complex from a scalar field defined over a tetrahedral mesh using a multi-stage, region-growing approach. We also simplify a Morse–Smale complex through the cancellation of critical point (zero-gradient-point) pairs. There are two types of cancellations—saddle point–extrema cancellations and 1-saddle–2-saddle cancellations. We have developed rules for determining when a cancellation is valid and rules for reconnecting a simplified complex.

A multi-resolution representation of a Morse–Smale complex can be obtained through a simplification hierarchy. We have defined rules for valid simplifications of a complex and developed a hierarchy based on independence of cancellations. We declare two cancellations to be independent when their affected areas in the Morse–Smale complex do not overlap. We also represent a hierarchy as a directed acyclic graph that encodes the independence of cancellations. A multi-resolution reconstruction of the complex can be obtained by cutting the hierarchy graph.

Additional work has been done under this subcontract for error estimation in the context of rendering large multi-resolution data sets. We have developed an algorithm to estimate the error associated with changing resolution levels in a data set stored as a multi-resolution hierarchy. Our method utilizes a pre-processing step that calculates intermediate resolution levels based on error to allow for interactive modification of transfer functions used in a volume rendering application. Our approach balances storage overhead cost and quality of error estimates in support of an efficient method for choosing the subset of a volumetric data set that will be rendered in a near-optimal way.

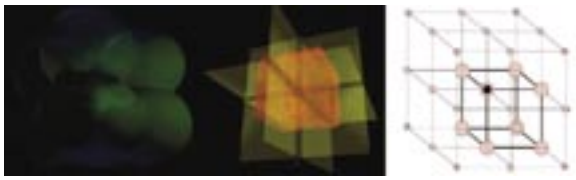

Iso-surface of test data set, stable and unstable manifolds, and associated Morse–Smale complex.

A saddle-maximum cancellation (left) and a 1-saddle–2-saddle cancellation.

Element Agglomeration AMGe for Contact Problems

Principal Investigators

Ana H. Iontcheva and Randolph E. Bank, University of California, San Diego

Collaborator

Panayot S. Vassilevski, CASC

The goal of this project was to develop algebraic multigrid methods for solving constrained-minimization problems, mainly contact problems in linear elasticity, discretized on general unstructured meshes, using the finite element method, in particular

- Signorini's problem
- Two-body contact problem
- Obstacle problem

For the solution of the Signorini's problem—contact of a linearly elastic body with a rigid frictionless foundation—we have developed two multilevel algorithms

- Multilevel subspace minimization algorithm
- FAS-constrained optimization algorithm

These schemes utilize element agglomeration coarsening away from the contact boundary, which allows for a straightforward construction of coarse-level approximations that automatically satisfy the fine-grid constraints. The two algorithms provide monotone reduction of the energy functional throughout the multilevel cycle. A code (in C++) has been developed and a paper has been written for the solution of the Signorini's problem. The results from this paper have been extended to 3D.

For the solution of the two-body contact problem, contact of two linearly elastic bodies in a mortar-method-based algorithm is the target of our future research and development. It will utilize the already developed methods for a single-body Signorini's contact problem.

Hypre for Symmetric Generalized Eigenvalue Problems

Principal Investigator

Andrew Knyazev, University of Colorado, Denver

Collaborator

Charles Tong, CASC

Symmetric eigenvalue problems are crucially important in structured mechanics and electronic structure calculations. The goal of the project is to implement a modern eigenvalue solver into Hypre in order to take advantage of Hypre high-quality preconditioners for parallel clusters with a user-friendly interface. The Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) method for symmetric eigenvalue problems has been chosen for implementation as one of the most promising eigensolvers in the current literature. The LOBPCG apparently combines the fast convergence of the block's Lanczos method with the simplicity and robustness of the block's steepest descent and allows the preconditioner to be plugged directly into the eigensolver without any inner-outer iteration schemes. The result is a well-written native Hypre C implementation of LOBPCG that works in three Hypre interfaces and is almost as efficient as the Hypre PCG linear solver.

The basic functionality of the code has been tested on University of Colorado, Denver and LLNL clusters on a number of test problems, including non-self-consistent electronic structure calculations eigenproblems, and the code is checked into the Hypre Alpha. Presently, this is the first and only preconditioned eigenvalue solver for parallel clusters.

An entirely new implementation of the LOBPCG for eigenvalue problems has been written as a part of the Hypre, consistent with Hypre Krylov-based linear solvers. The distinctive features of the new LOBPCG code are as follows.

1) The code is rewritten from scratch in order to implement the **lopbcg.c** code at the same level of abstraction as the Hypre **pcg.c** code.
2) The LOBPCG now works for three Hypre interfaces—IJ, Struct and SStruct—and includes test drivers for every interface. As of August 31, 2004, the LOBPCG drivers reflect the latest changes in the corresponding Hypre drivers and are prepared to serve as combined drivers for both linear systems and eigenvalue problems.
3) A new Hypre type, hypre_MultiVector, is introduced, and a preliminary implementation of Hypre MultiVector functions based on the existing parallel vector types (hypre_ParVector, hypre_StructVector and hypre_SStructVector) is used in the LOBPCG code. A significant future effort by the Hypre team is necessary to turn it into an actual MultiVector with efficient implementation of Hypre MultiVector functions by eliminating redundant MPI calls and by using Basic Linear Algebra Subprograms (BLAS).
4) The code solves generalized eigenvalue problems, as well as regular eigenvalue problems.
5) The code allows the use of constraints.

Preconditioning of Finite-Element Saddle-Point Problems

Principal Investigators

Raytcho D. Lazarov and Joseph E. Pasciak, Texas A&M University

Collaborator

Panayot Vassilevski, CASC

The goal of this research is the efficient solution of the discrete equations that result from finite-element or finite-difference approximation of partial-differential equations of mathematical physics. This involves the development and analysis of algorithms especially tailored for execution on medium- to large-scale parallel computing platforms. Our approach involves the application of theoretical tools from the analysis of partial differential equations to motivate and analyze new computational algorithms.

For implementation of highly accurate methods and efficient algorithms, we need approximation methods that provide greater flexibility in the grid generation process, increase the portability of various approximation methods and computer implementations, enhance the capabilities of coarsening strategy in parallel algebraic multi-grid methods, and provide a natural and practical way for parallel domain decomposition methods and parallel adaptive methods based on a posteriori error analysis. Our investigations produce competitive algorithms that can be used in various codes for complex applications in physics and engineering.

Main results of the research

1. A new, inexact Newton algorithm for the solution of second-order problems with higher-order nonlinearities was proposed and analyzed. This approach was based on a stability analysis in Sobolev spaces of order greater than one.
2. A new multigrid algorithm was proposed and analyzed for an electromagnetic problem. This method involves strengthening a curl-curl term by adding a discrete "grad-div" term. The resulting two-level multigrid algorithm was analyzed.
3. A stabilization framework for the discontinuous Galerkin finite-element method was developed. It was implemented for Raviart–Crouzeix non-conforming finite elements and tested on three-dimensional problems of elasticity.
4. A new stable scheme of exponential fitting type was proposed and studied for convection-diffusion problems.

Biological Fluid Flow in Micro-Electro-Mechanical Systems

Principal Investigator

Dorian Liepmann, University of California, Berkeley

Collaborator

David Trebotich, CASC

Our investigation aims to characterize the flow behavior of biological macromolecule solutions in silicon microfluidic devices. Digital Particle Image Velocimetry (DPIV) is used to quantify the velocity fields of DNA-laden solution flow under microfluidic conditions. Deviations from Newtonian flow fields brought about by the viscoelastic fluid rheology, concentration effects and conformational changes of the molecules will be assessed for flows through a variety of microfluidic geometries, such as straight channels, contractions, and expansions. Experimental results will be used to validate a computational design tool for bio-detection microdevices being developed at LLNL.

The characterization of flows containing high concentrations of macromolecules is critical for the optimal design of microfluidic systems for biochemical analyses. Since the fluid flow behavior at the microscale differs from the macroscale, a simple scaling down of processes may be insufficient to elicit the same performance. This thereby motivates research directed toward the characterization of the microflow behavior of the biological macromolecule DNA, commonly used in biochemical assays. Studies have shown that λ-DNA can undergo dramatic stretching in microfluidic flow with conformational changes occurring in both elongational and shear flows (Shrewsbury et al. 2001). The device geometry and flow rate, as well as the viscosity, concentration and relaxation time of the solution, influence the conformation of the DNA.

A 1.67:1 microfabricated abrupt planar contraction design was chosen as a test device (see figure). Experimental pressure drops for 450 μg/ml DNA (semi-dilute) solutions across abrupt contraction were measured. Velocity fields for flows of Newtonian and semi-dilute DNA solutions through abrupt contraction and abrupt expansion (inverted contraction device) were quantified using DPIV. No flow regions were observed near the contraction corners at slower flow rates and there is visual evidence of recirculation at faster flow rates.

Validation of the pressure testing system was conducted by measuring water flow through a narrow tube and comparing to theoretical solution. Pressure drop data was collected for flows of Newtonian (water) and non-Newtonian (DNA-laden) solutions through gradual and abrupt contraction devices. The effects of DNA solution concentration were explored and the onset of elastic, non-Newtonian behavior determined. DPIV was used to quantify flow fields for water flow in straight rectangular channels. A comparison study was conducted to determine the minimum interrogation region size for accurate data processing.

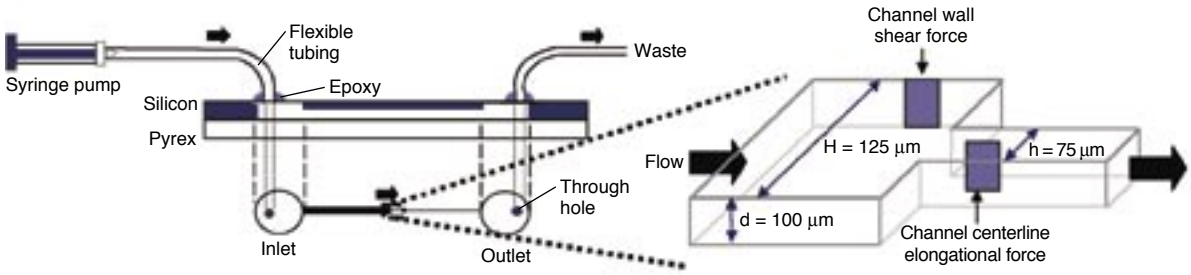

A microfabricated abrupt planar contraction design test device.at a ratio of 1.67 : 1.

Enabling Large-Scale Data Access

Principal Investigator **Ling Liu,** Georgia Institute of Technology

Collaborator **Terence Critchlow,** CASC

One of the ultimate goals of the Enabling Large Scale Data Access project is to produce a fully automated, end-to-end wrapper code generator through the design, development, and integration of service class descriptions with the XWRAP systems. The main idea is to provide mechanisms that enable XWRAP systems to take a generic description of a class of search interfaces (service classes), as well as the URL of a particular interface that is a member of the given class, to produce a functional wrapper that will take a class-specific query and produce XML views of the query results obtained through this particular interface. In the second year of this project, we focused on three main efforts.

(1) Extend the Web Spider developed under the previous subcontract to generate a set of Java 1.4 wrappers for interfaces that it identifies as matching a service-class description (SCD). Collectively, these wrappers will query all valid combinations of user input as specified by the SCD and the parameters enumerated by the interface. Each wrapper will be able to handle multiple types of errors (e.g., timeouts, connection refused, page not found) and return an appropriate message to the calling routine. The new version of the Web Spider is called DynaBot.

(2) In the DynaBot development, we have updated the Web-crawling component of the spider to be consistent with the interface identification component, updated the httpUnit library, and resolved JavaScript errors that prevent it from being used against some popular bio-portal sites, such as the National Center for Biotechnology Information's Basic Local Alignment Search Tool (BLAST) interface.

(3) We have tested the spider's ability to interact with the interface identification component by crawling the Web, starting at the dbCAT Web page and comparing the interfaces it identifies and the wrappers it generates with those identified and generated manually. This includes identifying and removing all obvious problems with the previous implementation of the spider, especially those identified in the second-year effort.

DYNABot System Architecture

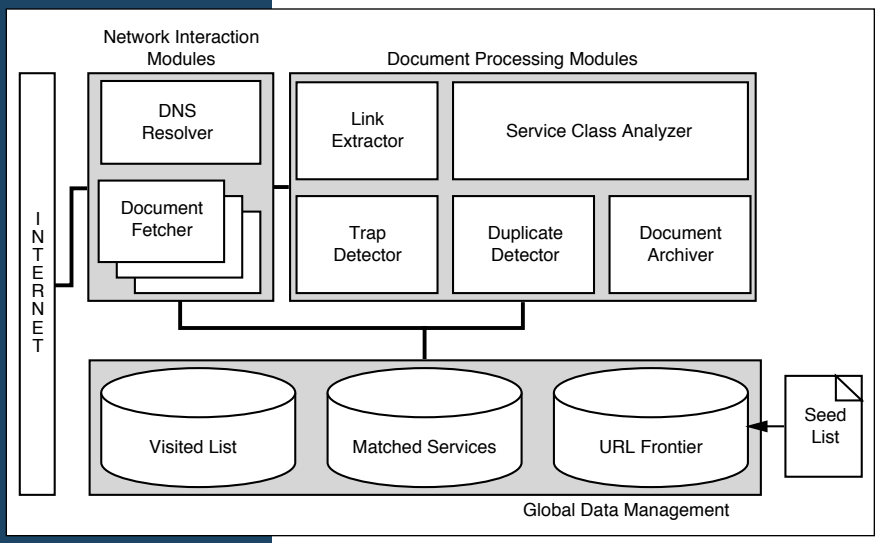

The components that make up a crawler are divided among three major component groups: network interaction modules, global storage and associated data managers, and document processing modules. The simplest crawlers require mechanisms for retrieving documents and determining if a particular URL has been seen. More advanced crawlers will include features like mirror-site detection and trap-avoidance algorithms. DynaBot utilizes an advanced crawler architecture for source discovery and adds a document processor that can determine if a dynamic Web source is related to a particular domain of interest. The architecture of Dynabot is shown at left.

For a more detailed description of this research, see the full report at: **http://disl.cc.gatech.edu/LDRD/**

Large Graph Visualization

Principal Investigator

Kwan-Liu Ma, University of California, Davis

Collaborator

Marvin Christensen, NAI

The Computer Incident Advisory Capability (CIAC) currently collects a great deal of data from sensors placed on a handful of computer networks. These sensors record activity on the networks so that the recorded information can be analyzed to provide useful information about both normal and anomalous activity on the networks.

This project focuses on one of the key significant technical challenges for interpreting this data. That challenge is a task in data mining—given the huge pool of data provided by the sensors, how can the interesting features be identified and classified? We address the problem by constructing a visualization system implementing novel methods of data analysis and information display. The process of exploring the space of a large data set using visualization is known as "visual data mining." It takes advantage of the pattern recognition facilities of the human visual system to detect patterns and anomalies in visual representations of abstract data.

We have developed a prototype system for visualizing TCP port data. The system permits analysts to discover the presence of any network security event that causes significant changes in the activity on ports. Since we currently only have access to very high-level data, the system is a very high-level tool and is useful mostly for uncovering high-level security events. Security events that consist of small details—an intrusion that includes only a few connections, for instance—are unlikely to be caught using this prototype system. Furthermore, since we have only obtained counts of activities rather than records of the activities themselves, the analysis can only go so far. The system can help identify suspicious traffic patterns, but it cannot see the traffic that caused the patterns. This is still useful, however, because analysts using this system can send the suspicious traffic signatures to analysts that have access to the full set of network traffic logs. The figure displays the current system interface.

Even in settings where only generalized information is available concerning network activity, we found many types of malicious activity can still be discovered using visualization. We have developed a tool that takes general, summarized network data and presents multiple, meaningful perspectives of the data and have demonstrated that this visualization leads to useful insights concerning network activity. Port scans of several types have been successfully detected, and many suspicious traffic patterns on individual ports have been uncovered. In addition, useful information about overall network traffic has been revealed, such as the rhythm of the traffic on commonly used ports as time progresses and the relationships between the various metrics used to describe port activity.

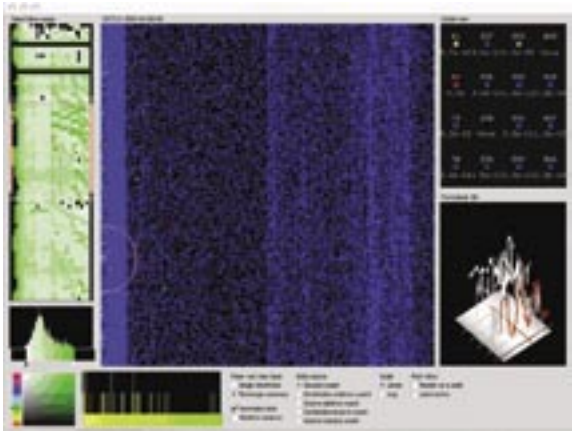

User interface of the TCP port data visualization system. The interface supports a drill-down process by allowing the analyst to explore high-level to lower-level information, from left to right, respectively.

Multi-Resolution Interactive Rendering of Large Scientific Data Sets Using an Image Cache

Principal Investigator **Kwan-Liu Ma,** University of California, Davis

Collaborator **Valerio Pascucci,** CASC

Large-scale simulations running on parallel supercomputers can produce very high-resolution information of the model phenomena, giving scientists the power to examine the phenomena at greater fidelity. Visualization enables scientists to better understand large amounts of data from such simulations. The problem with visualization is that conventional methods fail to handle the large amount of data. The data can take up too much system memory, as well as a considerable amount of time to render. Scientists are in need of tools that interactively search for and discover regions of interest within their simulations. Waiting for a long batch-mode visualization is not an option. In this project, we aimed at developing an interactive visualization solution enabling scientists to freely and effectively explore their simulation data.

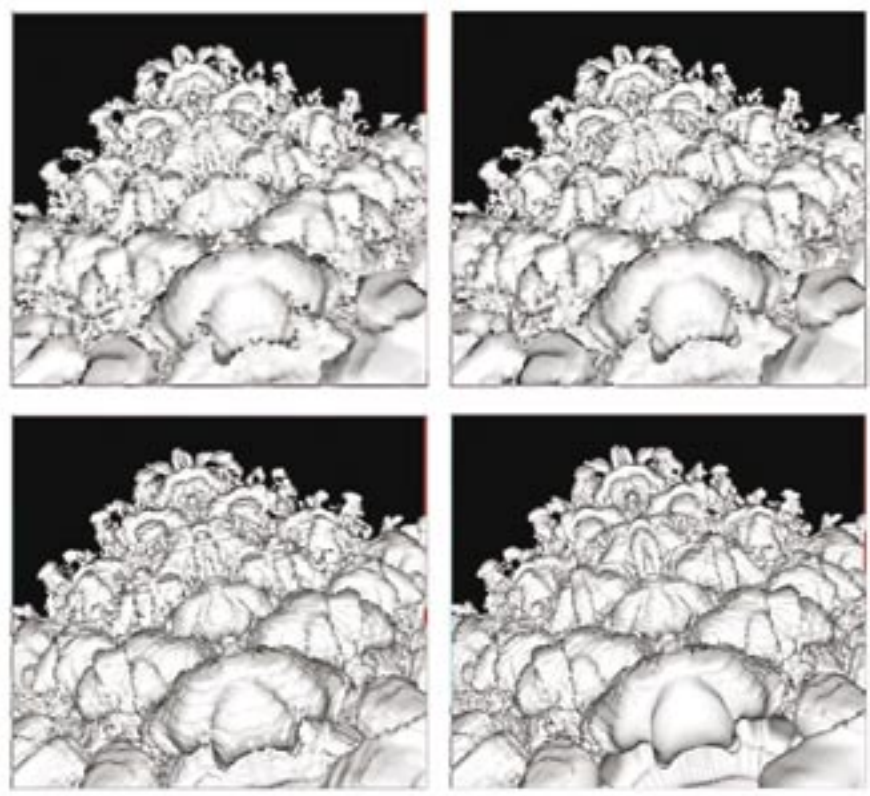

An example of the same scene rendered at different data resolutions. It is clear to see that as data resolution increases, images quality is greatly improved.

This project began as an attempt to provide a system that is capable of rendering large multiresolution data sets while taking advantage of a cluster computer. Our work involved the extension of a multi-layered image cache system (MLIC) developed at LLNL in several ways. First, the rendering engine was modified to access data sets stored in the Visualizations Schemes for Ultimate Scalability (ViSUS) cache oblivious to multi-resolution representation. This has improved the performance of the rendering engine with minimal loss in accuracy. Second, we removed the reliance on the Visualization Toolkit (VTK), which has greatly simplified the system implementation and its performance. Next, the system now runs on LLNL's new cluster computers, making it possible to use high-resolution output devices like a large PowerWall display. Finally, the user interface was revamped to be more intuitive and easy to use. The new system gives an easy method of navigating through a large dataset while simultaneously creating high-quality visualizations.

The extended MLIC system is a capable multi-resolution rendering system. It not only retains but enhances system interactivity, which allows a user to navigate through a dataset while it is being rendered. Rendering can be stopped and settings can be modified before completion, allowing a scientist to make changes on the fly. The new rendering classes have made the system render independent, which allows for the use of almost any rendering algorithm or rendering software.

As seen in the figure, the system can render a very coarse-resolution image, as well as a clear high-resolution image and several levels in between for the large Richtmyer-Meshkov turbulence simulation. This allows a scientist to have an idea about what is going to be rendered and possibly make changes. This is made possible by decoupling the rendering and displaying of images, allowing the system to always remain responsive even as the rendering is taking place. Responsiveness leads to easy navigation.

Open Source Software Technology for Transforming Scientific Problems

Principal Investigator

John Mellor-Crummey, Rice University

Collaborator

Dan Quinlan, CASC

Two leading software systems for source-to-source transformation and optimization of scientific programs are the ROSE infrastructure being developed at LLNL and the Open64/SL infrastructure being developed at Rice University. These efforts are complementary. The Open64/SL infrastructure has focused on software support for parsing and transforming Fortran 90-based programming models. ROSE has focused on software support for parsing, analysis and transformations of C++ based programming models. This project will build an open-source software technology to bridge these two infrastructures to allow ROSE's transformation capabilities to be applied to Fortran 90-based programs.

The project began with a study of Open64's WHIRL infrastructure, a study of LLNL's ROSE infrastructure (including its SAGE IR) and the construction of test programs to familiarize ourselves with ROSE. Currently, we have begun constructing a framework for the translation between WHIRL and ROSE. We have begun adding a C++ namespace to the ROSE code so that it can interoperate with the Open64 infrastructure and WHIRL. The translator itself is being constructed by modifying code that performs a WHIRL-to-WHIRL transformation to compile Co-array Fortran programs.

A Perl script will orchestrate the WHIRL-to-ROSE translation. This script will invoke mfef90 (the Open64 front end for Fortran 90) on a Fortran 90 source file to generate a file that contains the WHIRL representation source. The WHIRL external representation will be read by the translator proper (under construction), which will walk the WHIRL abstract syntax tree and translate each WHIRL operator into its corresponding operator in the SAGE IR. Once in the SAGE IR, ROSE-based tools can be applied to transform the resulting code. Next, the process will be reversed. The SAGE internal representation will be traversed and WHIRL will be reconstructed. The reconstructed WHIRL will be written out in a file using its external representation. Finally, the driver script will invoke Open64's whirl2ftool to regenerate Fortran from WHIRL.

We have been using Los Alamos National Laboratory's Parallel Ocean Program (POP), the National Academy of Sciences (NAS) parallel benchmarks, and the Adjoint MIT Ocean General Circulation Model to exercise Open64's support for analysis and transformation of Fortran. Currently, we are working with Open64's infrastructure on the parsing and regeneration of the POP code. With the exception of a minor problem unparsing arrays of structures, Open64 regenerates legal Fortran 90 code for POP.

Over the last several months, we have resolved a half dozen issues in Open64's support for Fortran, including correcting the very-high level WHIRL intermediate representation produced by the Fortran 90 parser, correcting unparsing of formatted I/O, correcting unparsing of array initializations. Currently, work is under way to correct problems with the constant pool produced in the WHIRL intermediate form and to produce a general infrastructure for correctly unparsing hierarchies of structures and arrays of structures.

Once we resolve known correctness issues, we will begin assembling a suite of codes for regression testing of Open64.

A Tightly Coupled Particle–Fluid Model for DNA-Laden Flows in Complex Microscale Geometries

Principal Investigators

Gregory H. Miller, University of California, Davis
P. Colella, D. T. Graves, D. F. Martin, P. O. Schwartz, LBNL

Collaborator

David P. Trebotich, CASC

We present a stable and convergent method for the computation of flows of DNA-laden fluids in microchannels with complex geometry. The numerical strategy combines a ball-rod model representation for polymers tightly coupled with a projection method for incompressible viscous flow. We use Cartesian-grid embedded boundary methods to discretize the fluid equations in the presence of complex domain boundaries. A sample calculation is presented showing flow through a packed array microchannel in 2D.

Approximate DNA in an actual microfluidic device is used for extraction. The left boundary condition is plug flow with a velocity of 0.1 cm/s; the right boundary is outflow (homogeneous Neumann); the top and bottom boundaries and the interior circular boundaries are a solid wall. The polymer is a 26-node approximation of DNA introduced near the left boundary as an inclined linear array after the fluid-flow field reached steady state. The polymer's trajectory causes it to become wrapped around the first circular element, where it remains pinned until the stochastic perturbations work it loose.

The fluid dynamic steps of this method are subject to an advective Courant–Friedrichs–Lewy stability condition only. The particle steps, without constraints, are also stable with this CFL timestep. When particles move far from the constraint manifold, however, the Lagrange multiplier algorithm of Ciccotti, et al. may diverge. We have found that the maximum particle displacement per time step for which the Ciccotti, et al. algorithm is stable may be extended for most systems by centering the constraint force at the conclusion of the timestep instead of the beginning.

We use a backward Euler time-stepping strategy that is formally first order accurate. To make the overall method second-order, it will be necessary to replace at least this with a Runge–Kutta time-stepping strategy.

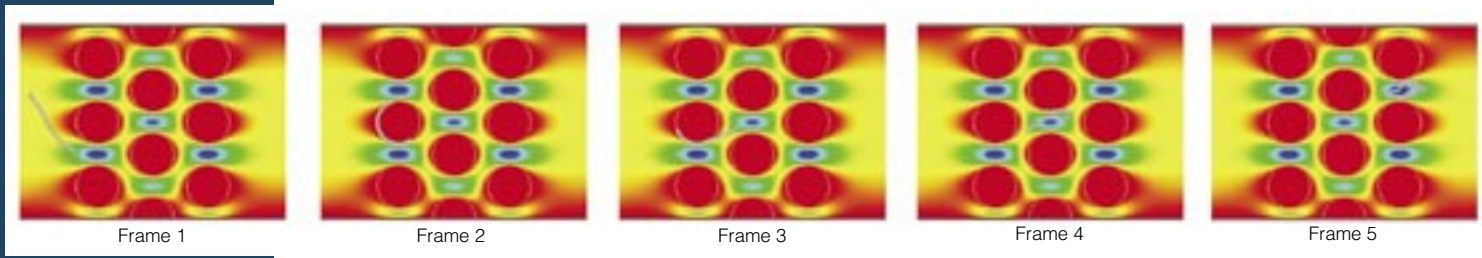

The time sequence of genomic DNA flowing in a 2D model of a packed-bed reactor PCR chamber. The DNA molecule enters from the left (frame 1), then wraps around a bead (frame 2), is loosened by hydrodynamic and Brownian forces (frame 3) and is swept out of the chamber by the flow field (frames 4 and 5). The color map indicates underlying flow field.

Clustered and Hierarchical Networks

Principal Investigator **Alex Pothen,** Old Dominion University

Collaborator **Edmond Chow,** CASC

We surveyed clustered and hierarchical representations of large networks available in the literature and made progress on two aspects of clustered graphs during the period of this contract.

First, we designed a new algorithm for discovering clusters in large networks that pays attention to the **local density** of the network. The density of a subgraph is the ratio of the number of edges joining nodes in the network to the maximum number of such edges possible (i.e., a complete subgraph or clique). In a large network, some subnetworks have high density, while others have low density, and a clustering algorithm that takes into account the local density would be able to do a better job of clustering.

In order to describe our algorithm, we need the concept of a ***k*-core**, a maximal subgraph in which every vertex is incident on at least k edges that belong to the subgraph (i.e., the degree of each vertex in the subgraph is at least k, when only edges in the subgraph are considered). The core number of a vertex is the maximum value of k as it belongs to a k-core, and the corresponding core is the maximum core to which the vertex belongs. The new clustering algorithm weights each vertex with the density of the connected component of the maximum core to which it belongs. The clustering proceeds by choosing a vertex with the highest weight to be the "seed vertex" for a cluster and then successively adds vertices whose weights are not too different from it to the cluster. After one cluster is complete, we begin with a new cluster in the same fashion.

This clustering algorithm has two advantages over existing ones. First, the weight of a vertex depends not on its degree but on the density of its neighborhood, thus clustering low-degree vertices together with high-degree vertices to which they are adjacent. Second, the algorithm can be efficiently implemented in time $O(K\,|E|)$, where K is the maximum core number of any vertex, and $|E|$ is the number of edges, in the network. We have implemented the weight computation phase of the algorithm and are currently working on the clustering phase.

Second, we are implementing a clustering algorithm for bipartite networks in which vertices in both parts need to be clustered. The algorithm relies on a singular vector of the vertex edge incidence matrix of the network. We followed the suggestion of Ravi Kannan, Santosh Vemula, and colleagues to speed up the singular-vector computation by computing with a submatrix of the incidence matrix. The columns in the submatrix are chosen with a probability proportional to the column weight. Then, the computed singular vector can be shown to be close to the singular vector of the larger matrix with high probability. Our preliminary results on the original incidence matrix shows that this approach works well, especially when combined with a combinatorial "local-smoothing" algorithm.

Spectral AMGe, ALE3D & FOSPACK

Principal Investigator

John W. Ruge, Front Range Scientific Computations, Boulder, CO

Collaborator

Rob Falgout, CASC

The majority of the work performed in FY 2004 has been on FOSPACK, a package for the discretization and solution of First-Order System Least-Squares (FOSLS) formulations of a range of PDEs. Algebraic Multigrid (AMG) is incorporated as the linear equation solver, and some of the effort has also gone into extending the range and improving the efficiency of that method. Progress has been made in several of the proposed areas of study and is summarized below.

AMG for Constraints

Slide-surface constraints were added to a two-dimensional, two-body elasticity problem. The constraints specified that two faces in contact remained in contact under application of a force, coupling the normal displacements along the two faces (and allowing for slippage in the tangential direction). Tests were conducted incorporating the constraints either strongly by essentially eliminating rows/columns from the matrix, or weakly by including them as weighted penalty terms in the functional to be minimized. Both resulted in an asymptotic V-cycle convergence factors that were nearly identical to those obtained without the constraints (where the two bodies were fully decoupled).

Parallel FOSPACK

Initial parallelization of 3D FOSPACK and coupling to Hypre/BoomerAMG has been completed for linear PDE systems with weak boundary conditions. The code has been ported to a 128-processor Beowulf cluster, and improvements in implementation resulted in better parallel performance. Treatment of nonlinear problems and coupled nonlinear multi-physics systems has also been improved. Tests performed on a coupled nonlinear fluid-elastic model and Maxwell's equations problem have resulted in solutions consistent with those obtained from the scalar FOSPACK code. The major effort for 3D FOSPACK is the introduction of higher-order discretizations, which are necessary for obtaining reasonable accuracy for 3D problems while keeping storage requirements within practical limits.

AMG Algorithm Enhancements

High-operator complexity can be a problem for AMG, especially in large 3D applications. In such cases, the coarse grid operators can be too large or too dense. Aggressive coarsening and multi-pass interpolation have been implemented in both the serial and parallel codes, resulting in markedly lower complexity, although some degradation of convergence is noted. A number of alternatives are being tested in the scalar code in an effort to both reduce complexity while retaining (or even improving) convergence.

Several variations of the basic AMG algorithm are being studied in an effort to improve its robustness. The use of compatible relaxation (CR) for choosing the coarser grids, combined with interpolation designed to minimize the trace of the coarse grid operator, looks quite promising. For some problems, the standard AMG interpolation is not accurate enough, particularly those with near-null space components. For this, an adaptive approach, in which interpolation is explicitly modified to better match the smoothest (slowest-to-converge) components, is being tested. The current approach is a hybrid of the recently developed adaptive Smoothed Aggregation (aSA) and a more straightforward adaptive AMG (aAMG).

Development of Efficient and Robust Algorithms for the Numerical Solution of Partial Differential Equations for High-Speed Reactive and Nonreactive Flow

Principal Investigator **Donald W. Schwendeman,** Rensselaer Polytechnic Institute

Collaborator **William Henshaw,** CASC

The work carried out under this subcontract involved the development and use of an adaptive numerical method for the accurate calculation of high-speed reactive and nonreactive flows on overlapping grids. In the reactive case, the flow is modeled by the reactive Euler equations with various choices for the equation of state and reaction rate model.

A numerical method has been developed to solve the nonlinear hyperbolic partial differential equations in the model. The method uses an unsplit, shock-capturing scheme, a Godunov-type scheme to compute fluxes, and a Runge–Kutta error control scheme to compute the source term modeling the chemical reactions. Two approximate Riemann solvers are now available and can be used for a general (non-ideal) equation of state. An exact Riemann solver is also available for an ideal, gamma-law equation of state.

Adaptive mesh refinement (AMR) has been implemented in order to locally increase grid resolution. The numerical method may also be applied to nonreactive flow problems, in which case the reactive source term is simply set to zero. The numerical method uses composite overlapping grids to handle complex flow geometries in either two or three dimensions. The code is part of the Overture–OverBlown framework of object-oriented codes, and the development has occurred in close collaboration with Bill Henshaw and Lori Diachin, and other members of the Overture group within the Center for Applied Scientific Computing.

During the period of this subcontract, a number of tasks were accomplished, including

- Further development of the numerical method for 2D moving meshes
- Implementation of additional Riemann solvers, including a Harten–Lax–van Leer (HLL) approximate Riemann solver for a general equation of state and an exact Riemann solver for an ideal equation of state
- An extension of the numerical method to handle axisymmetric flow
- An extension of the numerical method to handle multi-material, non-reactive flow